


Attention Please: Deep Dive into Image Captioning Model Efficiency

Mithun “Mit” Ramesh

This project investigated the performance of transformer-based models in image captioning tasks, critically examining the limitations of generalized models like VGG19, ResNet, and InceptionX. We created a Transformer seq2seq model from scratch in order to replicate our own image caption prediction model that was trained on the Flickr8K Dataset. In creating this self-created model, we experimented with many hyperparameters and experimental decisions (Different inference algorithms, using RNN/LSTM, etc). We then compared this performance to the performance of un-tuned pre-trained models with a fine-tuned model on specific datasets.. Our objectives were to evaluate the state-of-the-art in image captioning and assess the adaptability of these models to specialized datasets, a crucial aspect for applications like medical imaging. This was a group project, and for my portion of the capstone project, I attempted to expand the original Flickr8k dataset to include the Flickr30k dataset and the COCO120k dataset; the preprocessing classes can be found in our github repository. To improve the training efficiency, I implemented multithreading on both the Flickr8k and Flickr30k datasets. Additionally, I developed a tokenizer and embedding class to process the textual data in the datasets for the from-scratch model. To evaluate the performance of our seq2seq model, I also implemented the Greedy Search method as an evaluation method.

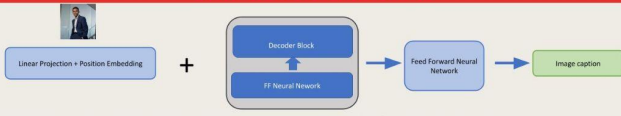
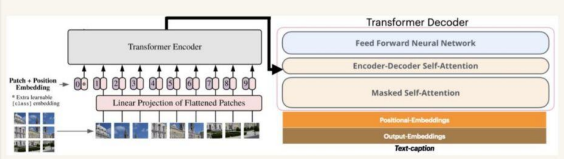
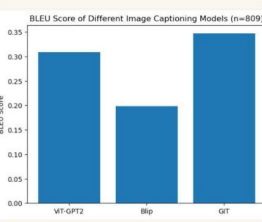

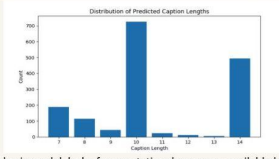
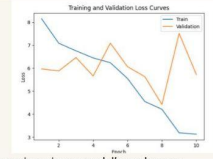


Attention Please! Deep-dive into Image Captioning Efficacy

Kevin Kang, Ethan Polley, Mithun Ramesh, Siddharth Somasi

CSCI

1430

Motivation	Problem	Goal
<p>The success of seminal research paper “Attention is All You Need” has highlighted the power of self-attention mechanisms in language modeling tasks, and has spurred a wave of research into transformer-based models. As companies like OpenAI continue to push the boundaries of what’s possible with generative text models, it’s clear that the race is far from over.</p> <p>Image captioning is a challenging task in computer vision that requires a deep understanding of both visual and linguistic content. Our group aims to investigate and compare the performance of different transformer-based models for image captioning.</p>	<p style="text-align: center;">Foundational Models: Too Big to Fail?</p> <p>While image extractors like VGG19, ResNet, and InceptionX have achieved impressive results on a wide range of image recognition tasks, <u>there is a growing concern that their generalized nature may lead to blind spots in the AI</u>. This lack of specificity poses a significant problem for image captioning, as it can limit the model’s ability to accurately describe and contextualize images. Our group aim to investigate this issue by comparing efficacy of un-tuned base image captioning models relative to a fine-tuned image captioning model on a specific dataset.</p>	<p style="text-align: center;">Our Action Items</p> <ol style="list-style-type: none"> 1. Implement a simple image captioning model by hand that utilized seq2seq architecture. Use different datasets and model architectures 2. Compare the efficacy of pre-existing models such as opensource ViT-GPT2, Microsoft Research Team’s GIT, and Salesforce’s BLP 3. Fine-tune the GIT base models, retraining on niche datasets, to investigate model efficacy prior and after fine-tuning
<div style="display: flex; justify-content: space-between;"> <div style="width: 48%;"> <p style="text-align: center; background-color: #f00; color: white; padding: 2px;">Basic model architecture</p>  <ul style="list-style-type: none"> The basic model was fairly straightforward: it utilized an image-feature extractor (experimented with VGG16, ResNet, InceptionV3) and Spacy-processed captions and was fed to a Seq2Seq model that predicted the next token (only one that predicted sensible tokens). We utilized GLOVE embeddings Experimented with RNN/LSTM/GRU and a TransformerBlock architecture on smaller subset of Flickr8K Dataset. Experimented with Greedy Search, which seemed to have limitations, and implemented Beam Search Had Flickr30k and COCO120k data created, but could not find resources to test. Possibly would improve performance </div> <div style="width: 48%;"> <p style="text-align: center; background-color: #f00; color: white; padding: 2px;">ViT-GPT2 Model Architecture</p>  </div> </div>		
<p style="text-align: center; background-color: #f00; color: white; padding: 2px;">Results</p> <div style="display: flex; align-items: center;">  <div style="margin-left: 20px;"> <p style="font-size: 0.8em;">‘a stuffed animal with a flower in it’</p> <p style="font-size: 0.8em;">↓</p> <p style="font-size: 0.8em;">‘a green and red toy with red eyes’</p>  </div> </div>	<p style="text-align: center; background-color: #f00; color: white; padding: 2px;">More results and Limitations</p> <div style="display: flex; justify-content: space-between;">   </div> <ul style="list-style-type: none"> For the basic model, lack of computational resources available became a large issue in our model’s weaknesses Best BLEU score: 0.05. Not very strong, but we believe the use of Greedy Search + low data limits affected our performance. Greedy Search favored very specific tokens and required careful engineering. We implemented Beam Search but it wasn’t computationally efficient. Performance on fine-tuned obviously did not go as plan given lack of generality on general data 	
<p style="text-align: center; font-size: 0.8em;">References</p> <p style="font-size: 0.7em;">Vaswani, Ashish, et al. “Attention is All You Need.” <i>Advances in Neural Information Processing Systems</i>, vol. 30, 2017</p> <p style="font-size: 0.7em;">Lambada. “Pokemon-Blp-Captions.” <i>Huggingface.co</i>, 13 Dec. 2022</p> <p style="font-size: 0.7em;">adityaj105. “Flickr 8k Datasets.” <i>Kaggle.com</i>, 2020. www.kaggle.com/datasets/adityaj105/flickr8k</p>	<p style="text-align: center; font-size: 0.8em;">Acknowledgements</p> <p style="font-size: 0.7em;">As senior capstoners, we would like to thank all those — whether Professors, fellow students, or close friends — that have helped us along the way to the finish line!</p>	